

WHAT IS CLAIMED IS:

1 1. A method for managing multiple resources in a system including at least
2 one host, network, and a storage space comprised of at least one storage system that each
3 host is capable of accessing over the network, comprising:

4 measuring and monitoring a plurality of service level parameters indicating a state
5 of the resources in the system;

6 determining values for the service level parameters;

7 determining whether the service level parameter values satisfy predetermined
8 service level thresholds;

9 indicating whether the service level parameter values satisfy the predetermined
10 service thresholds; and

11 determining a modification of one at least one resource deployment or
12 configuration if the value for the service level parameter for the resource does not satisfy
13 the predetermined service level thresholds.

1 2. The method of claim 1, wherein the monitored service level parameter
2 comprises one of a performance parameter and an availability level of at least one system
3 resource.

1 3. The method of claim 2, wherein the service level performance parameters
2 that are monitored are members of a set of performance parameters comprising: a
3 downtime during which the at least one application is unable to access the storage space;
4 a number of times the at least one application host was unable to access the storage
5 space; a throughput in terms of bytes per second transferred between the at least one host
6 and the storage; and an I/O transaction rate.

1 4. The method of claim 1, wherein the modification of resource deployment
2 comprises at least one of adding additional instances of the resource and modifying a
3 configuration of the resource.

PCT/US2014/045001

1 5. The method of claim 1, wherein a time period is associated with one of the
2 monitored service parameters, further comprising:

3 determining a time during which the value of the service level parameter
4 associated with the time period does not satisfy the predetermined service level threshold;
5 and

6 generating a message indicating that the determined time exceeds the time period
7 if the determined time exceeds the time period associated with the monitored service
8 parameter.

1 6. The method of claim 5, wherein a customer contracts with a service
2 provider to provide the system at agreed upon service level parameters, further
3 comprising:

4 transmitting a service message to the service provider after determining that the
5 value of the service level parameter does not satisfy the predetermined service level; and
6 transmitting the message indicating failure of the value of the service level
7 parameter for the time period to both the customer and the service provider.

1 7. The method of claim 1, further comprising writing to a log information
2 indicating whether the service level parameter values satisfy the predetermined service
3 thresholds.

1 8. The method of claim 1, wherein determining the modification of the at
2 least one resource deployment further comprises:

3 analyzing the resource deployment to determine at least one resource that
4 contributes to the failure of the service level parameter values to satisfy the threshold;
5 determining whether any additional instances of the determined at least one
6 resource that contributes to the failure of the service level parameter is available; and
7 allocating at least one additional instance of the determined at least one resource
8 to the system.

1 9. The method of claim 8, wherein analyzing the resource deployment
2 comprises performing a bottleneck analysis.

1 10. The method of claim 8, further comprising:
2 determining characteristics of access to the resources by applications operating at
3 the service level;
4 if there are no additional instances of the determined at least one resource, then
5 determining whether the access characteristics exceed predetermined access
6 characteristics; and
7 indicating that the service level is not sufficient due to a change in the access
8 characteristics.

1 11. The method of claim 10, wherein the access characteristics include
2 read/write ratio, Input/Output (I/O) size, and percentage of access being either sequential
3 or random.

1 12. The method of claim 10, wherein the predetermined access characteristics
2 are specified in a service level agreement that indicates the thresholds for the service
3 level parameter values.

1 13. The method of claim 1, wherein a plurality of applications at different
2 service levels are accessing the resources in the system, wherein requests from
3 applications using a higher priority service receive higher priority than requests from
4 applications operating at a lower priority service, wherein determining the modification
5 of the at least one resource deployment further comprises:
6 increasing the priority associated with the service level whose service level
7 parameter values fail to satisfy the predetermined service level thresholds.

SEARCHED INDEXED
SERIALIZED FILED

1 14. The method of claim 13, wherein determining the modification of the at
2 least one resource deployment further comprises:

3 analyzing the resource deployment to determine at least one resource that
4 contributes to the failure of the service level parameter values to satisfy the thresholds;
5 determining whether any additional instances of the determined at least one
6 resource that contributes to the failure of the service level parameter is available; and
7 allocating at least one additional instance of the determined at least one resource
8 to the system, wherein the priority is increased if there are no additional instances of the
9 at least one resource that contributes to the failure.

1 15. The method of claim 1, wherein one service level parameter value
2 indicates a time throughput of Input/Output operations between the at least one host and
3 the storage space has been below a throughput threshold, and wherein determining the
4 additional resource allocation further comprises determining at least one of host adaptor,
5 network, and storage resources to add to the configuration.

1 16. The method of claim 1, further comprising:
2 invoking an operation to implement the determined additional resource allocation.

1 17. The method of claim 1, wherein the service level parameters specify a
2 predetermined redundancy of resources, further comprising:
3 detecting a failure of one component;
4 determining whether the component failure causes the resource deployment to fall
5 below the predetermined redundancy fo resources; and
6 indicating whether the component failure causes the resource deployment to fall
7 below the predetermined redundancy threshold.

PENTEX INFORMATION SYSTEMS INC.

1 18. A system for managing multiple resources in a system including at least
2 one host, network, and a storage space comprised of at least one storage system that each
3 host is capable of accessing over the network, comprising:

4 means for measuring and monitoring a plurality of service level parameters
5 indicating a state of the resources in the system;

6 means for determining values for the service level parameters;
7 means for determining whether the service level parameter values satisfy
8 predetermined service level thresholds;

9 means for indicating whether the service level parameter values satisfy the
10 predetermined service thresholds; and

11 means for determining a modification of at least one resource deployment or
12 configuration if the value for the service level parameter for the resource does not satisfy
13 the predetermined service level thresholds.

1 19. The system of claim 18, wherein the service level performance parameters
2 that are monitored are members of a set of performance parameters comprising: a
3 downtime during which the at least one application is unable to access the storage space;
4 a number of times the at least one application was unable to access the storage space; a
5 throughput in terms of bytes per second transferred between the at least one application
6 and the storage; and an I/O transaction rate.

1 20. The system of claim 18, wherein the modification of resource deployment
2 comprises at least one of adding additional instances of the resource and modifying a
3 configuration of the resource.

1 21. The system of claim 18, wherein a time period is associated with one of
2 the monitored service parameters, further comprising:

PATENT
MAILED
01/12/2018

3 means for determining a time during which the value of the service level
4 parameter associated with the time period does not satisfy the predetermined service level
5 threshold; and
6 means for generating a message indicating that the determined time exceeds the
7 time period if the determined time exceeds the time period associated with the monitored
8 service parameter.

1 22. The system of claim 18, wherein the means for determining the
2 modification of the at least one resource deployment further performs:
3 analyzing the resource deployment to determine at least one resource that
4 contributes to the failure of the service level parameter values to satisfy the threshold;
5 determining whether any additional instances of the determined at least one
6 resource that contributes to the failure of the service level parameter is available; and
7 allocating at least one additional instance of the determined at least one resource
8 to the system.

1 23. The system of claim 22, further comprising:
2 means for determining characteristics of access to the resources by applications
3 operating at the service level;
4 means for determining whether the access characteristics exceed predetermined
5 access characteristics if there are no additional instances of the determined at least one
6 resource; and
7 means for indicating that the service level is not sufficient due to a change in the
8 access characteristics.

1 24. The system of claim 18, wherein a plurality of applications at different
2 service levels are accessing the resources in the system, wherein requests from
3 applications using a higher priority service receive higher priority than requests from

4 applications using a lower priority service, wherein determining the modification of the
5 at least one resource deployment further comprises:
6 increasing the priority associated with the service level whose service level
7 parameter values fail to satisfy the predetermined service level thresholds.

1 25. A system for managing multiple resources in a system including at least
2 one host, network, and a storage space comprised of at least one storage system that each
3 host is capable of accessing over the network, comprising:
4 a processing unit;
5 a computer readable medium accessible to the processing unit;
6 program code embedded in the computer readable medium executed by the
7 processing unit to perform:
8 (i) measuring and monitoring a plurality of service level parameters
9 indicating a state of the resources in the system;
10 (ii) determining values for the service level parameters;
11 (iii) determining whether the service level parameter values satisfy
12 predetermined service level thresholds;
13 (iv) indicating whether the service level parameter values satisfy the
14 predetermined service thresholds; and
15 (v) determining a modification of at least one resource deployment or
16 configuration if the value for the service level parameter for the resource does not
17 satisfy the predetermined service level thresholds.

1 26. The system of claim 25, wherein the service level performance parameters
2 that are monitored are members of a set of performance parameters comprising: a
3 downtime during which the at least one application is unable to access the storage space;
4 a number of times the at least one application was unable to access the storage space; a
5 throughput in terms of bytes per second transferred between the at least one application
6 and the storage; and an I/O transaction rate.

1 27. The system of claim 25, wherein the program code for determining the
2 modification of the resource deployment comprises at least one of adding additional
3 instances of the resource and modifying a configuration of the resource.

1 28. The system of claim 25, wherein a time period is associated with one of
2 the monitored service parameters, wherein the program code is further executed by the
3 processing unit to perform:

4 determining a time during which the value of the service level parameter
5 associated with the time period does not satisfy the predetermined service level threshold;
6 and

7 generating a message indicating that the determined time exceeds the time period
8 if the determined time exceeds the time period associated with the monitored service
9 parameter.

10

1 29. The system of claim 25, wherein the program code for determining the
2 modification of the at least one resource deployment further causes the processing unit to
3 perform:

4 analyzing the resource deployment to determine at least one resource that
5 contributes to the failure of the service level parameter values to satisfy the threshold;

6 determining whether any additional instances of the determined at least one
7 resource that contributes to the failure of the service level parameter is available; and

8 allocating at least one additional instance of the determined at least one resource
9 to the system.

1 30. The system of claim 29, wherein the program code is further executed by
2 the processing unit to perform:

3 determining characteristics of access to the resources by applications operating at
4 the service level;

5 determining whether the access characteristics exceed predetermined access
6 characteristics if there are no additional instances of the determined at least one resource;
7 and
8 indicating that the service level is not sufficient due to a change in the access
9 characteristics.

1 31. The system of claim 25, wherein a plurality of applications at different
2 service levels are accessing the resources in the system, wherein requests from
3 applications using a higher priority service receive higher priority than requests from
4 applications using a lower priority service, wherein the program code for determining the
5 modification of the at least one resource deployment further causes the processing unit to
6 perform:

7 increasing the priority associated with the service level whose service level
8 parameter values fail to satisfy the predetermined service level thresholds.

1 32. An article of manufacture including code for managing multiple resources
2 in a system including at least one host, network, and a storage space comprised of at least
3 one storage system that each host is capable of accessing over the network, wherein the
4 code is capable of causing operations comprising:

5 measuring and monitoring a plurality of service level parameters indicating a state
6 of the resources in the system;

7 determining values for the service level parameters;

8 determining whether the service level parameter values satisfy predetermined
9 service level thresholds;

10 indicating whether the service level parameter values satisfy the predetermined
11 service thresholds; and

12 determining a modification of one at least one resource deployment or
13 configuration if the value for the service level parameter for the resource does not satisfy
14 the predetermined service level thresholds.

PCT/US2014/061004

1 33. The article of manufacture of claim 32, wherein the monitored service
2 level parameter comprises one of a performance parameter and an availability level of at
3 least one system resource.

1 34. The article of manufacture of claim 33, wherein the service level
2 performance parameters that are monitored are members of a set of performance
3 parameters comprising: a downtime during which the at least one host is unable to access
4 the storage space; a number of times the at least one host was unable to access the storage
5 space; a throughput in terms of bytes per second transferred between the at least one host
6 and the storage; and an I/O transaction rate.

1 35. The article of manufacture of claim 32, wherein the modification of
2 resource deployment comprises at least one of adding additional instances of the resource
3 and modifying a configuration of the resource.

1 36. The article of manufacture of claim 32, wherein a time period is associated
2 with one of the monitored service parameters, further comprising:
3 determining a time during which the value of the service level parameter
4 associated with the time period does not satisfy the predetermined service level threshold;
5 and
6 generating a message indicating that the determined time exceeds the time period
7 if the determined time exceeds the time period associated with the monitored service
8 parameter.
9

1 37. The article of manufacture of claim 36, wherein a customer contracts with
2 a service provider to provide the system at agreed upon service level parameters, further
3 comprising:
4 transmitting a service message to the service provider after determining that the
5 value of the service level parameter does not satisfy the predetermined service level; and

6 transmitting the message indicating failure of the value of the service level
7 parameter for the time period to both the customer and the service provider.

1 38. The article of manufacture of claim 32, further comprising writing to a log
2 information indicating whether the service level parameter values satisfy the
3 predetermined service thresholds.

1 39. The article of manufacture of claim 32, wherein determining the
2 modification of the at least one resource deployment further comprises:
3 analyzing the resource deployment to determine at least one resource that
4 contributes to the failure of the service level parameter values to satisfy the threshold;
5 determining whether any additional instances of the determined at least one
6 resource that contributes to the failure of the service level parameter is available; and
7 allocating at least one additional instance of the determined at least one resource
8 to the system.

1 40. The article of manufacture of claim 39, wherein analyzing the resource
2 deployment comprises performing a bottleneck analysis.

1 41. The article of manufacture of claim 39, further comprising:
2 determining characteristics of access to the resources by applications operating at
3 the service level;
4 if there are no additional instances of the determined at least one resource, then
5 determining whether the access characteristics exceed predetermined access
6 characteristics; and
7 indicating that the service level is not sufficient due to a change in the access
8 characteristics.

1 42. The article of manufacture of claim 41, wherein the access characteristics
2 include read/write ratio, Input/Output (I/O) size, and a percentage of access being either
3 sequential or random.

1 43. The article of manufacture of claim 41, wherein the predetermined access
2 characteristics are specified in a service level agreement that indicates the thresholds for
3 the service level parameter values.

1 44. The article of manufacture of claim 32, wherein a plurality of applications
2 at different service levels are accessing the resources in the system, wherein requests
3 from applications using a higher priority service receive higher priority than requests
4 from applications operating at a lower priority service, wherein determining the
5 modification of the at least one resource deployment further comprises:
6 increasing the priority associated with the service level whose service level
7 parameter values fail to satisfy the predetermined service level thresholds.

1 45. The article of manufacture of claim 44, wherein determining the
2 modification of the at least one resource deployment further comprises:
3 analyzing the resource deployment to determine at least one resource that
4 contributes to the failure of the service level parameter values to satisfy the thresholds;
5 determining whether any additional instances of the determined at least one
6 resource that contributes to the failure of the service level parameter is available; and
7 allocating at least one additional instance of the determined at least one resource
8 to the system, wherein the priority is increased if there are no additional instances of the
9 at least one resource that contributes to the failure.

1 46. The article of manufacture of claim 32, wherein one service level
2 parameter value indicates a time throughput of Input/Output operations between the at
3 least one host and the storage space has been below a throughput threshold, and wherein

4 determining the additional resource allocation further comprises determining at least one
5 of host adaptor, network, and storage resources to add to the configuration.

1 47. The article of manufacture of claim 32, further comprising:
2 invoking an operation to implement the determined additional resource allocation.

1 48. The article of manufacture of claim 32, wherein the service level
2 parameters specify a predetermined redundancy of resources, further comprising:
3 detecting a failure of one component;
4 determining whether the component failure causes the resource deployment to fall
5 below the predetermined redundancy fo resources; and
6 indicating whether the component failure causes the resource deployment to fall
7 below the predetermined redundancy threshold.